

FortiOS™ Handbook - Transparent Mode

VERSION 5.4.1



FORTINET DOCUMENT LIBRARY

<http://docs.fortinet.com>

FORTINET VIDEO GUIDE

<http://video.fortinet.com>

FORTINET BLOG

<https://blog.fortinet.com>

CUSTOMER SERVICE & SUPPORT

<https://support.fortinet.com>

<http://cookbook.fortinet.com/how-to-work-with-fortinet-support/>

FORTIGATE COOKBOOK

<http://cookbook.fortinet.com>

FORTINET TRAINING SERVICES

<http://www.fortinet.com/training>

FORTIGUARD CENTER

<http://www.fortiguard.com>

FORTICAST

<http://forticast.fortinet.com>

CLI REFERENCE

<http://cli.fortinet.com>

END USER LICENSE AGREEMENT

<http://www.fortinet.com/doc/legal/EULA.pdf>

FEEDBACK

Email: techdocs@fortinet.com



Wednesday, December 13, 2017

FortiOS™ Handbook - Transparent Mode

TABLE OF CONTENTS

Change Log	5
Introduction	6
Introduction to Transparent Mode	7
What is Transparent Mode?	7
Transparent Mode Features	7
Transparent Mode Installation	9
Installing a FortiGate in Transparent mode	9
Using a Virtual Wire Pair to Simplify Transparent Mode	10
Management IP configuration	11
In-band management details and example	11
Out-of-band management details and example	12
Networking in Transparent Mode	13
Packet Forwarding	13
MAC learning and L2 Forwarding Table	13
Broadcast, Multicast, and Unicast Forwarding	14
Multicast Processing	15
Source MAC Addresses	16
ARP Table	17
Verifying the Forwarding Database	17
Spanning Tree BPDUs Forwarding	17
Non-IPv4 Ethernet frames forwarding	18
Network Address Translation (NAT)	18
Configuring SNAT	18
Configuring DNAT	19
VLANs and Forwarding Domains	20
VLANs in Transparent Mode	20
Forwarding Domains in Transparent Mode	20
VLANs vs Forwarding Domains	21
VLAN Forwarding	21
Unknown VLAN Processing	22
VLAN Trunking and MAC Address Learning	22
VLAN Translation	22
Inter-VDOM links between NAT/Route and Transparent VDOMs	22
Replay Traffic Scenario	23

Packet Forwarding using Cisco Protocols.....	23
Configuration Example.....	25
Firewalls and Security in Transparent Mode.....	27
Firewall Policy Look Up.....	27
Firewall Session List.....	27
Security Scanning.....	28
Integrated Switch Fabric Access Control List Policies.....	28
IPsec VPN in Transparent Mode.....	30
Using IPsec VPNs in Transparent Mode.....	30
Example 1: Remote Sites with Different Subnets.....	31
Example 2: Remote Sites on the Same Subnet.....	37
Using FortiManager and FortiAnalyzer.....	44
High Availability in Transparent Mode.....	45
Virtual Clustering.....	45
MAC Address Assignment.....	46
Best Practices.....	47

Change Log

Date	Change Description
June 24, 2016	New publication.
Dec 13, 2017	Updated "VLAN Forwarding" on page 21 and "Using FortiManager and FortiAnalyzer" on page 44.

Introduction

This guide explains how to use a FortiGate in Transparent Mode. It includes the following sections:

- [Introduction to Transparent Mode](#): an overview of Transparent mode, including available features.
- [Transparent Mode Installation](#): instructions for installing a FortiGate in Transparent mode.
- [Networking in Transparent Mode](#): how networking is configured in Transparent mode.
- [Firewalls and Security in Transparent Mode](#): information about using firewalls and security scanning in Transparent mode.
- [IPsec VPN in Transparent Mode](#): configuring IPsec VPNs using FortiGates in Transparent mode.
- [Using FortiManager and FortiAnalyzer](#): using external management and logging.
- [High Availability in Transparent Mode](#): configuring Transparent FortiGates in HA mode.
- [Best Practices](#): the best practices for using Transparent mode.

Introduction to Transparent Mode

This section contains a [basic overview of Transparent mode](#), as well as a [chart showing features available in Transparent mode](#).

What is Transparent Mode?

A FortiGate unit can operate in one of two modes: Transparent or NAT/Route mode.

In Transparent mode, the FortiGate is installed between the internal network and the router. In this mode, the FortiGate does not make any changes to IP addresses and only applies security scanning to traffic. When a FortiGate is added to a network in Transparent mode, no network changes are required, except to provide the FortiGate with a management IP address. Transparent mode is used primarily when there is a need to increase network protection but changing the configuration of the network itself is impractical.

In NAT/Route mode, a FortiGate unit is installed as a gateway or router between two networks. This allows the FortiGate to hide the IP addresses of the private network using network address translation (NAT).

Transparent Mode Features

Different FortiOS features are available depending on whether your FortiGate is in Transparent or NAT/Route mode. The following table shows which features are available for each mode.



For a FortiGate in Transparent mode, the maximum number of Interfaces per VDOM is 254. This value includes both physical and virtual interfaces.

For any other maximum values, please consult the Maximum Values Table, available at docs.fortinet.com.

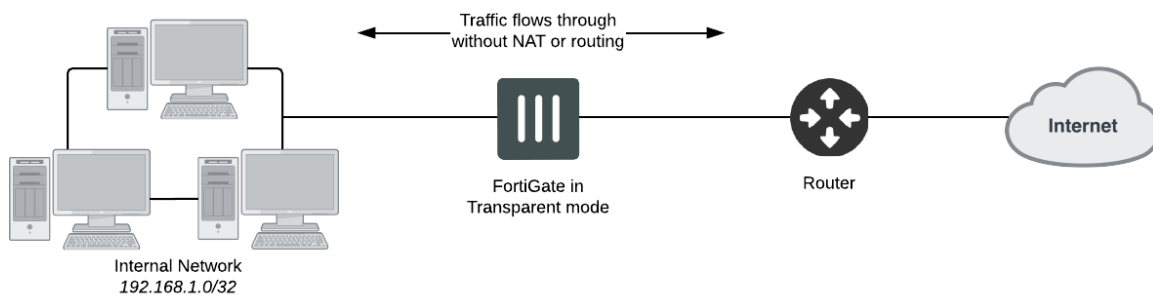
Feature/capability	NAT	Transparent	Comment
Unicast Routing / Policy Based routing	YES	NO	
VIP / IP pools / NAT	YES	YES	Configurable from CLI only in transparent mode
Multicast routing	YES	NO - Options available to forward multicast packets	

Feature/capability	NAT	Transparent	Comment
L2 forwarding	NO	YES	In Transparent mode, other frames than IP can be forwarded but with no UTM processing.
Firewall (packet filtering / NAT / Authentication)	YES	YES	
IPv6 capable	YES	YES	
Traffic shaping - TOS classification	YES	YES	
Hardware acceleration	YES	YES	
All Security profile features	YES	YES	
Cooperative Security Fabric	YES	YES	
FortiView	YES	YES	
IPsec gateway	YES	YES - Policy based mode only	
SSL gateway	YES	NO	
High-Availability (HA) - Virtual Cluster	YES	YES	
802.3ad (LACP / port aggregation)	YES	YES	
HA port redundancy	YES	YES	FortiGate hardware dependent
802.1q - VLAN trunking	YES	YES	
802.1d - Spanning Tree	NO	NO - option to forward BPDUs	
Logging and reporting (disk and memory logging, FortiCloud, syslog, and FortiAnalyzer)	YES	YES	
Managed by FortiManager	YES	YES	

Transparent Mode Installation

This section contains information about [installing a FortiGate in Transparent mode](#) and using a [virtual wire pair](#), to simplify Transparent mode. It also contains information about [configuring the FortiGate's management IP](#).

Installing a FortiGate in Transparent mode



Changing to Transparent mode removes most configuration changes made in NAT/Route mode. To keep your current NAT/Route mode configuration, backup the configuration using the **System Information** widget, found in the **Dashboard**.

1. Before connecting the FortiGate unit to your network, go to the **Dashboard** and locate enter the following command into the **CLI Console**:

```
config system settings
  set opmode transparent
  set manageip <address and netmask>
  set gateway <address>
end
```

2. Access the web-based manager by browsing to the new management IP.
3. (Optional) The FortiGate unit's DNS Settings are set to use FortiGuard DNS servers by default, which is sufficient for most networks. However, if you need to change the DNS servers, go to **Network > DNS** and add **Primary** and **Secondary** DNS servers. Select **Apply**.
4. If your network uses IPv4 addresses, go to **Policy & Objects > IPv4 Policy** and select **Create New** to add a security policy that allows users on the private network to access the Internet.

If your network uses IPv6 addresses, go to **Policy & Objects > IPv6 Policy** and select **Create New** to add a security policy that allows users on the private network to access the Internet. If the IPv6 menu option is not available, go to **System > Feature Select**, turn on **IPv6**, and select **Apply**. For more information on IPv6 networks, see the IPv6 Handbook.

Set the **Incoming Interface** to the internal interface and the **Outgoing Interface** to the Internet-facing interface (typically WAN1). You will also need to set **Source Address**, **Destination Address**, **Schedule**, and **Service** according to your network requirements. You can set these fields to the default all/ANY settings for now but should create the appropriate objects later after the policies have been verified.

5. Make sure the **Action** is set to **ACCEPT**. Select **OK**.



It is recommended to avoid using any security profiles, such as AntiVirus or web filtering, until after you have successfully installed the FortiGate unit. After the installation is verified, you can apply any required security profiles.

For more information about using security profiles, see the Security Profiles handbook.

6. Go to the **Dashboard** and locate the **System Resources** widget. Select **Shutdown** to power off the FortiGate unit.
Alternatively, you can also use the CLI command `execute shutdown`.
7. Connect the FortiGate unit between the internal network and the router.
8. Connect the Internet-facing interface to the router's internal interface and connect the internal network to the FortiGate using an internal port (typically port 1).
9. Power on the FortiGate unit. You will experience downtime before the FortiGate unit starts up completely.


Results

Users on the internal network are now able to browse to the Internet. They should also be able to connect to the Internet using any other protocol or connection method that you defined in the security policy.



If a FortiGate unit operating in Transparent mode is installed between your internet network and a server that is providing a network service to the internal network, such as DNS or DHCP, you must add a wan1-to-internal policy to allow the server's response to flow through the FortiGate unit and reach the internal network.

Using a Virtual Wire Pair to Simplify Transparent Mode



In FortiOS 5.4, virtual wire pair replaces the feature port pairing from earlier firmware versions. Unlike port pairing, virtual wire pair can be used for a FortiGate in either Transparent or NAT/Route mode.

A virtual wire pair consists of two interfaces that have no IP addresses and all traffic received by one interface in the pair can only be forwarded out the other; as controlled by firewall policies. Since the interfaces do not have IP addresses, you can insert a virtual wire pair into a network without having to make any changes to the network.



Interfaces used in a virtual wire pair cannot be used for admin access to the ISFW FortiGate. Before creating a virtual wire pair, make sure you have a different port configured to allow admin access using your preferred protocol.

1. Go to **Network > Interfaces** and select **Create New > Virtual Wire Pair**. Add two ports to the virtual wire pair. These ports cannot be part of a switch, such as the default **internal/lan** interface.
2. Go to **Policy & Objects > IPv4 Virtual Wire Pair Policy** and create a policy will allow traffic to flow between the two ports. Give the policy an appropriate **Name**. Select the direction that traffic is allowed to flow. Configure the other firewall options as needed.
3. If necessary, create a second virtual wire pair policy allowing traffic to flow between the ports in the opposite direction.

Traffic can now flow between the two ports. Go to **FortiView > Policies** to see traffic flowing through both policies.

Management IP configuration

A FortiGate in Transparent mode can be assigned with a single IP address for remote access management and multiple static routes can be configured. This can be used if in-band management wants to be applied.

When out-of-band management is desired (dedicated interface for remote management access), it is recommended to use a separate VDOM in NAT/Route mode.

In-band management details and example

The management IP address is bound to all ports or VLANs belonging to the same VDOM. Remote access services are subject to the same rules as in NAT/Route mode, and must be enabled/disabled on each port.

Example of management IP configuration in Transparent mode:

```
config system settings
    set manageip 10.1.1.100/255.255.255.0
end
config router static
    edit 1
        set gateway 10.1.1.254
    next
end
config system interface
    edit port1
        set allowaccess ping ssh https snmp
    end
```

It is also possible to add a second IP address for management and additional default routes:

```
config system settings
```

```
set opmode transparent
set manageip 192.168.182.136/255.255.254.0 10.1.1.1/255.255.255.0
end

config router static
edit 1
    set gateway 192.168.183.254
next
edit 2
    set gateway 10.1.1.254
next
end
```



`ping-server` (dead gateway detection) is not supported in Transparent mode.

Out-of-band management details and example

When VDOM is enabled and the VDOMs are operating in Transparent mode, it is recommended, to avoid L2 loops and allow more routing flexibility, to keep one VDOM (generally the root VDOM) in NAT/Route mode, with one or more VLAN or physical interface as out-of-band management.



The management VDOM must have IP connectivity to the Internet to allow communication with the FDS and retrieve services information (AV, IPS, Fortiguard filtering, Support...). All syslog and FortiManager communication also go through the management VDOM.

Networking in Transparent Mode

This section contains information about networking concepts in Transparent mode. It contains the following topics:

- [Packet Forwarding](#)
- [Network Address Translation \(NAT\)](#)
- [VLANs and Forwarding Domains](#)
- [Inter-VDOM links between NAT/Route and Transparent VDOMs](#)
- [Packet Forwarding using Cisco Protocols](#)
- [Configuration Example](#)

Packet Forwarding

The following sections include information about configuring packet forwarding in Transparent mode:

- [MAC learning and L2 Forwarding Table](#)
- [Broadcast, Multicast, and Unicast Forwarding](#)
- [Multicast Processing](#)
- [Source MAC Addresses](#)
- [ARP Table](#)
- [Verifying the Forwarding Database](#)
- [Spanning Tree BPDUs Forwarding](#)
- [Non-IPv4 Ethernet frames forwarding](#)

MAC learning and L2 Forwarding Table

When operating in Transparent mode, a FortiGate behaves like an L2 switch in accordance with 802.1d principles:

- The forwarding database (FDB) is populated with the network devices MAC addresses during a MAC learning process, based on the source addresses seen in the Ethernet frames ingressing a FortiGate port. Static MAC entries can also be configured using the following CLI command:

```
config system mac-address-table
  edit 00:01:02:03:04:05
    set interface "port3"
  next
end
```

The FDB table can be verified with the following command: `diagnose netlink
brctl name host TP.b`

- Ethernet IP frames forwarding is based on known MAC address on each port.
- As Spanning Tree is not running on the FortiGate, a port that comes up goes immediately into forwarding or flooding state. This last state will not occur once unicast MAC addresses are present in the FDB.



If the FortiGate in Transparent mode bridges traffic to a router or host using a virtual MAC for one direction and a different physical MAC for the other direction (for example when VRRP or HSRP protocols are used), it is highly recommended to create a static MAC entry for the virtual MAC. This is to make sure that the virtual MAC address is present in the FDB.

Broadcast, Multicast, and Unicast Forwarding

In Transparent mode, IPv4 packets are typically only forwarded by the FortiGate from a port to another port when a firewall policy is matched with action ACCEPT.

Below are exceptions.

- **L2 (IP) Broadcast frames forwarding:**



L2 (IP) means a L2 frame type 0x0800 (IP) or 0x0806 (ARP)

- **ARP:** by default, ARP broadcasts and ARP reply packets are flooded/forwarded on all ports or VLANs belonging to the same forwarding domain, without the need of firewall policies between the ports. This default behavior is necessary to allow the population of the FDB and allow further firewall policy lookup (see section Transparent mode Firewall processing for more details). This option is configurable at the interface settings level with the parameter `arpforward` (enabled by default).
- **Non-ARP:** To forward non-ARP broadcasts, the following CLI command is used:

```
config system interface
  edit "port2"
    set broadcast-forward enable
  next
end
```

- **L2 (IP) Multicast frames forwarding:** the FortiGate does not forward frames with multicast destination MAC addresses by default. Multicast traffic such as one used by routing protocols or streaming media may need to traverse the FortiGate which should not interfere this communication.

Fortinet recommends that the FortiGate is set up using Multicast policies. This allows for greater control and predictability on traffic behavior. However Multicast traffic may be forwarded through a Transparent mode device using the `multicast-skip-policy` setting. This is detailed in the section ["Multicast Processing" on page 15](#)

- **L2 (IP) Unicast frames forwarding:** a frame with a unicast destination MAC address is subject to firewall processing before being forwarded (see ["Firewall Policy Look Up" on page 27](#) for more details). This does not apply to ARP replies.

Multicast Processing

In Transparent mode, a FortiGate does not forward frames with multicast destination MAC addresses by default. If multicast traffic is required, multicast policies are recommended to allow finer control of this traffic.

Forwarding all multicast traffic with policy

Multicast traffic may have to be forwarded through a Transparent mode device using the `multicast-skip-policy` system setting. This is the configuration for this solution:

```
config system settings
    set multicast-skip-policy enable
end
```

In that case, no check is performed on sources/destinations/interfaces. A multicast packet received on an interface is flooded unconditionally to all interfaces (except the incoming interface) belonging to the same forwarding domain.

Configuring firewall multicast-policy

The use of `firewall multicast-policy` allows a finer control over the multicast packets. Hereafter are some commented examples. Note that the parameter `multicast-skip-policy` mentioned above must be left to disabled.

Those policies can only be configured from the CLI.

1- Simple policy

```
config firewall multicast-policy
    edit 1
        set action accept
    next
end
```

In that case, no check is performed on sources/destinations/interfaces. A multicast packet received on an interface is flooded unconditionally to all interfaces (except the incoming interface) belonging to the same forwarding domain.

2- To restrict incoming and outgoing interfaces:

```
config firewall multicast-policy
    edit 1
        set srcintf "port1"
        set dstintf "port2"
        set action accept
    next
end
```

3- To be more restrictive (example to allow RIP2 packets from port1 to port2 and sourced by 10.10.0.10):

```
config firewall multicast-policy
edit 1
set srcintf "port1"
set srcaddr 10.10.0.10 255.255.255.255
set dstintf "port2"
set dstaddr 224.0.0.9 255.255.255.255
set action accept
next
end
```

4- This policy will allow all 224.0.0.0/255 range (OSPF, RIPv2, DVMRP...) from port1 to port2

```
config firewall multicast-policy
edit 1
set srcintf "port1"
set dstintf "port2"
set dstaddr 224.0.0.0 255.255.255
set action accept
next
end
```

Source MAC Addresses

When a FortiGate is in Transparent Mode, it does not typically alter the original source and destination address of packets that flow through the unit. Because of this, end devices do not “see” the MAC address of the FortiGate. However, if network address translation (NAT) is enabled by a firewall policy, the source MAC address will be the MAC address of the FortiGate's management interface.

IP packets that are initiated by the FortiGate (remote management, access to FortiGuard server...) are sent in L2 Ethernet frames that have a source MAC address of the interface in the virtual domain (VDM) with the lowest MAC address. Below is an example with port2 and port3 in the same VDOM, remote access done via port2, but the sniffer trace showing MAC address of port2. The address of port2 is shown in bold.

```
diagnose hardware deviceinfo nic port2
[...]
```

```
Current_HWaddr      00:09:0F:85:3F:C4
Permanent_HWaddr    00:09:0F:85:3F:C4
```

```
fgt300 (global) # diagnose hardware deviceinfo nic port3
```

```
[...]
Current_HWaddr      00:09:0F:85:3F:C5
Permanent_HWaddr    00:09:0F:85:3F:C5
```

```
diagnose sniffer packet port3 "port 80" 6
3.774236 port3 -- 192.168.171.165.2619 -> 192.168.182.136.80: syn 3961770249
0x0000  0009 0f85 3fc4 0009 0f09 3204 0800 4500      ....? .... 2...E.
0x0010  0030 8071 4000 7e06 98d7 c0a8 aba5 c0a8      .0.q@.~ .....
0x0020  b688 0a3b 0050 ec23 d109 0000 0000 7002      ...;.P.# ..... p.
0x0030  ffff d7e7 0000 0204 05b4 0101 0402
```


ARP Table

In Transparent mode, the Address Resolution Protocol (ARP) table is used in the following situations:

- For IP traffic received or originated by the FortiGate itself, and in destination of the management device or next-hop.
- When IPsec is used, the FortiGate uses its ARP table to forward the traffic from the IPsec tunnel to the local destination host(s).

All other forwarding decision is based on the Forwarding Database (FDB) table or optional settings.

Verifying the Forwarding Database

To view all instances of the forwarding database (FDB), use the following CLI command:

```
diagnose netlink brctl list
```

Example

```
FGT # diagnose netlink brctl list

list bridge information
1. root.b      fdb: size=256    used=6    num=7    depth=2    simple=no
2. mgmt.b      fdb: size=256    used=5    num=4    depth=2    simple=no
Total 2 bridges
```

Here above we can see two bridge instances for 2 VDOMs in Transparent mode: `root` and `mgmt`.

- This command will dump the L2 forwarding table for each VDOM bridge instance:

```
diagnose netlink brctl name host <VDOM_name>.b
```

Example for the root VDOM:

```
FGT# diag netlink brctl name host root.b

show bridge control interface root.b host.
fdb: size=256, used=6, num=7, depth=2, simple=no
Bridge root.b host table
port no device devname mac addr          ttl      attributes
  2    7    wan2    02:09:0f:78:69:00      0        Local Static
  5    6    trunk_1 02:09:0f:78:69:01      0        Local Static
  3    8    dmz      02:09:0f:78:69:01      0        Local Static
  4    9    internal 02:09:0f:78:69:02      0        Local Static
  3    8    dmz      00:80:c8:39:87:5a      194
  4    9    internal 02:09:0f:78:67:68      8
  1    3    wan1     00:09:0f:78:69:fe      0        Local Static
```

Spanning Tree BPDUs Forwarding

Spanning tree Bridge Protocol Data Units (BPDUs) are not forwarded by default in Transparent Mode. To forward spanning tree BPDUs, the following setting can be applied on each interface where this is required:

```
config system interface
edit port1
set stpforward enable
```

```
next
end
```

Non-IPv4 Ethernet frames forwarding

In the situation where non IP frames (or non Ethernet II) frames need to be accepted on a port, the parameter `l2forward` can be enabled (disabled by default). This can be used to forward frames such as PPPoE PADI, Appletalk, on other ports belonging to the same forwarding domain.

The procedure is the following:

```
config system interface
edit port1
set l2forward enable
next
edit port2
set l2forward enable
next
end
```

Network Address Translation (NAT)

While NAT is generally not used by a FortiGate in Transparent mode, both [source network address translation \(SNAT\)](#) and [destination network address translation \(DNAT\)](#) can be configured.

Configuring SNAT

Source Network Address Translation (SNAT) is an option available in Transparent mode and configurable in CLI only, using the following commands:

```
config firewall ippool
edit "nat-out"
set endip 192.168.183.48
set startip 192.168.183.48
set interface vlan18_p3
next
end

config firewall policy
edit 3
set srcintf "vlan160_p2"
set dstintf "vlan18_p3"
set srcaddr "all"
set dstaddr "all"
set action accept
set ippool enable
set poolname "nat-out"
set schedule "always"
set service "ANY"
set nat enable
next
end
```

The sniffer trace below shows the source IP 192.168.182.93 being source translated to 192.168.183.48:

```
fgt300 (TP) # diagnose sniffer packet any "host 10.2.2.1" 4

interfaces=[any]
filters=[host 10.2.2.1]
4.891970 vlan160_p2 in 192.168.182.93 -> 10.2.2.1: icmp: echo request
4.892003 vlan18_p3 out 192.168.183.48 -> 10.2.2.1: icmp: echo request
4.892007 port3 out 192.168.183.48 -> 10.2.2.1: icmp: echo request
4.933216 vlan18_p3 in 10.2.2.1 -> 192.168.183.48: icmp: echo reply
4.933249 vlan160_p2 out 10.2.2.1 -> 192.168.182.93: icmp: echo reply
4.933253 port2 out 10.2.2.1 -> 192.168.182.93: icmp: echo reply
```

Configuring DNAT

The following example shows how to configure Destination Network Address Translation (DNAT) using a virtual IP on a FortiGate in Transparent Mode:

```
config firewall vip
  edit "vip1"
    set extip 192.168.183.48
    set extintf "vlan160_p2"
    set mappedip 192.168.182.78
  next
end

config firewall policy
  edit 4
    set srcintf "vlan160_p2"
    set dstintf "vlan18_p3"
    set srcaddr "all"
    set dstaddr "vip1"
    set action accept
    set schedule "always"
    set service "ANY"
  next
end
```



If the mappedip is on a different subnet than the management IP, the FortiGate must have a valid route to this destination

The sniffer trace below shows the destination IP 192.168.183.48 being translated to 192.168.182.78:

```
fgt300 (TP) # diagnose sniffer packet any "icmp" 4

interfaces=[any]
filters=[icmp]
4.126138 vlan160_p2 in 192.168.182.93 -> 192.168.183.48: icmp: echo request
4.126190 vlan18_p3 out 192.168.182.93 -> 192.168.182.78: icmp: echo request
4.126196 port3 out 192.168.182.93 -> 192.168.182.78: icmp: echo request
4.126628 vlan18_p3 in 192.168.182.78 -> 192.168.182.93: icmp: echo reply
4.126661 vlan160_p2 out 192.168.183.48 -> 192.168.182.93: icmp: echo reply
4.126667 port2 out 192.168.183.48 -> 192.168.182.93: icmp: echo reply
```

VLANs and Forwarding Domains

The following sections include information about configuring virtual local area networks (VLANs) and forwarding domains in Transparent mode:

- [VLANs in Transparent Mode](#)
- [Forwarding Domains in Transparent Mode](#)
- [VLANs vs Forwarding Domains](#)
- [VLAN Forwarding](#)
- [Unknown VLAN Processing](#)
- [VLAN Trunking and MAC Address Learning](#)
- [VLAN Translation](#)

VLANs in Transparent Mode

A VLAN configured on a physical port is used to classify a packet in a broadcast domain in ingress and to tag packet in egress. A VLAN on the FortiGate conforms to the standard 802.1q. The following rules apply to VLAN configuration:

- a VLAN ID can be used only once on the same physical port
- the same VLAN ID can be used on a different port
- the VLAN ID range is from 1 to 4094

Forwarding Domains in Transparent Mode

A forwarding domain is used to create separate broadcast domains and confine traffic across two or more ports. It also allows learning the same MAC in different VLANs (IVL).

A forwarding domain and its associated ID number are unique across one VDOM, or a FortiGate with VDOMs disabled. Each new VDOM will create a new bridge instance in the FortiGate.



Even though the forwarding domain ID is not in relation with the actual VLAN numbers, it is recommended, for maintenance and troubleshooting purposes, to configure one forwarding domain per VLAN and use the same forwarding domain ID as the VLANs ID.

Once forwarding domains are configured, it is possible to configure firewall policies only between ports or VLAN belonging to the same forwarding domain.

Figure 4 shows an example with three forwarding domains and VLANs configured. In this example, there are two VDOMs in Transparent Mode: root and MGMT. Forwarding domain 0 is the default on the FortiGate or VDOM in Transparent Mode.

- Root VDOM has:
 - 3 forwarding domains, 0, 340, and 341.
 - VLAN 340 configured on port1; packets will be tagged with ID 340
 - VLAN 341 configured on port1; packets will be tagged with ID 341
 - All other ports are untagged
- MGMT VDOM has got only the default forwarding domain 0

The expected behavior is the following:

- Packets untagged ingressing port1, port3 and port4 belong to the same broadcast domain in the root VDOM
- Packets tagged with VLAN 340 ingressing port1 and Packets untagged ingressing port2 belong to the same broadcast domain in the root VDOM
- Packets tagged with VLAN 341 ingressing port1 and Packets untagged ingressing port5 belong to the same broadcast domain in the root VDOM
- Packets untagged ingressing port6 belong to a different broadcast domain in the MGMT VDOM

Configuration example for forwarding domain 340:

```
config system interface
  edit "VLAN340"
    set forward-domain 340
    set interface "port1"
    set vlanid 340
  next
  edit "port3"
    set forward-domain 340
  next
end
```

VLANs vs Forwarding Domains

There are several differences between VLAN and a forwarding domain configured on a FortiGate in Transparent Mode:

- A forwarding domain is used to create separated broadcast domains between VLANs and allow independent VLAN learning - IVL (MAC addresses in the FDB). This would be equivalent to creating VLANs on a regular L2 switch.



When VLANs are used in the network, configuring different forwarding domains is essential to avoid broadcast duplications. See also section Default VLAN forwarding behavior for additional information.

- VLANs configured on interfaces are only used for tagging packets egressing the port and classifying packets at ingress.
- The packets processed by the direct interface (or port) itself are always sent untagged and must be received untagged.

VLAN Forwarding

VLAN forwarding allows you to forward all VLANs traffic of a trunk that was connecting two network devices and where the FortiGate has been introduced, without having to perform any further configuration.

It is recommended to configure forwarding domains for each VLAN and disable this parameter in order to avoid packet from looping into the trunk from one VLAN to another. By default, the parameter `vlanforward` is disabled on each physical interface of a FortiGate or VDOM in Transparent mode.

Unknown VLAN Processing

When a FortiGate receives a tagged frame with an unknown VLAN ID, the processing is the following:

- By default, the ports are set with the parameter `vlanforward enable`; in this scenario, all tagged frames are forwarded to the ports belonging to the same forwarding domain. This allows you to insert the FortiGate between two devices using trunk ports without any further configuration.
- When using forwarding domains in association with VLANs, `vlanforward` should be disabled whenever applicable. In this scenario, any new tagged frames with unknown VLAN IDs are dropped by the FortiGate. This is the solution recommended by Fortinet.

For an example VLAN configuration, see "Configuration Example" on page 25.

VLAN Trunking and MAC Address Learning

A FortiGate port becomes a trunk when 2 or more VLANs are configured on this port, in the same or different forwarding domains.



When trunks are configured on a FortiGate, it is essential to create forwarding domains, in order to avoid packets looping back on the VLANs of the trunk. This will confine all broadcasts and multicast traffic between the interfaces belonging to a same forward domain.

In the case where a trunk port is configured with a VLAN in a different forwarding domains, the MAC address of the network device connected to this port learns the FDB of each forwarding domain. This is Independent VLAN Learning (IVL).

VLAN Translation

The same forwarding domain can include several different VLANs. Therefore, a frame ingressing an interface with a certain VLAN ID can be forwarded to another port with another VLAN ID. This is sometimes referred as VLAN translation.

Inter-VDOM links between NAT/Route and Transparent VDOMs

Inter-VDOM links between NAT/Route and Transparent mode VDOMS can be useful for configurations where the NAT/Route VDOMs that share a common Internet service route, which can be routed through a Transparent VDOM that provides additional functionality, like common Security inspection, WAN optimization, explicit proxying and so on.

Other examples include:

- Performing SSL offloading in the Transparent mode VDOM and providing Internet access through a NAT/Route mode VDOM.

- Applying WAN optimization in a Transparent mode VDOM and other security features in the NAT/Route mode VDOM.
- Using a dedicated Transparent mode VDOM for the explicit web proxy in front of a NAT/Route mode VDOM that applies other security features.
- An ISP configuration with multiple per-tenant NAT/Route mode VDOMs all sharing a single Internet connection but where the ISP only presents a single routed subnet. Each tenant can then be assigned an IP from the subnet for their respective VDOM link interface while using a single physical port to connect to the ISP router.

For more information about inter-VDOM links, please refer to the Virtual Domains handbook.

Replay Traffic Scenario

Situations can arise where an identical TCP packet enters twice the FortiGate via 2 different ports. This can be due to a firewall or other network device redirecting packets out on the same port it has received it.

The FortiGate will in this condition detect a replay packet and drop it.

If the network topology or culprit devices cannot be changed to avoid this, the workaround on the FortiGate can be to disable TCP replay verification packets.

```
config system global
    set anti-replay | loose | strict | disable |
end
```

The debug flow diagnosis output hereafter shows the message indicating this condition:

```
id=20085 trace_id=179 msg="vd-VDOM_VLAN1 received a packet(proto=6, 10.10.253.9:10709
>10.10.248.5:25) from TO_EXTERNAL ."
id=20085 trace_id=179 msg="Find an existing session, id-00041475, original direction"
id=20085 trace_id=179 msg="replay packet, drop"
```

For additional diagnosis and troubleshooting procedures, go to <http://kb.fortinet.com>.

Packet Forwarding using Cisco Protocols

In order to pass Cisco Discover Protocol (CDP) or Cisco VLAN Trunk Protocol (VTP) packets through a FortiGate in Transparent mode, the parameter `stpforward` must be applied on the port configuration. VTP and CDP packets are sent to the destination MAC address 01-00-0C-CC-CC-CC.



A Cisco NATIVE VLAN carries CDP/VTP frames. The frames of this VLAN must be received on the FortiGate physical interfaces (not VLAN sub-interface). Physical interfaces are the only ones that can send/accept non-tagged packets.

The example below will allow CDP and VTP packets to be sent from port3 up to the Remote unit, through two VDOMs, via one physical port and three port aggregations.

Port and Port aggregation configuration:

```
config system interface
    edit "port1"
```

```
        set vdom "VD1"
    next
    edit "port2"
        set vdom "VD1"
    next
    edit "port3"
        set vdom "VD1"
        set stpforward enable
    next
    edit "port5"
        set vdom "VD3"
    next
    edit "port6"
        set vdom "VD3"
    next
    edit "port17"
        set vdom "VD2"
    next
    edit "port18"
        set vdom "VD2"
    next
    edit "port19"
        set vdom "VD2"
    next
    edit "port20"
        set vdom "VD2"
    next
    edit "LACP_VD2_IN"
        set vdom "VD2"
        set stpforward enable
        set type aggregate
        set member "port17" "port18"
    next
    edit "LACP_VD2_OUT"
        set vdom "VD2"
        set stpforward enable
        set type aggregate
        set member "port19" "port20"
    next
    edit "LACP_VD1"
        set vdom "VD1"
        set stpforward enable
        set type aggregate
        set member "port1" "port2"
    next
end
```



When using aggregation, the `stpforward` setting needs to be applied only on the port aggregation level, not on the physical port

This will also forward regular Spanning Tree BPDUs

Verification with a sniffer trace:

```
FGT# diagnose sniffer packet any "" 4
```



```

41.365434 port3 in llc unnumbered, ui, flags [command], length 72
41.365437 LACP_VD1 out llc unnumbered, ui, flags [command], length 72
41.365439 port2 out llc unnumbered, ui, flags [command], length 72
41.365479 LACP_VD2_IN in llc unnumbered, ui, flags [command], length 72
41.365482 LACP_VD2_OUT out llc unnumbered, ui, flags [command], length 72
41.365484 port19 out llc unnumbered, ui, flags [command], length 72

```


See above the CDP packet flow from port3, LACP_VD1 (port2), LACP_VD2_IN, LACP_VD2_OUT (port19).



The following sniffer trace command will filter only CDP or VTP packets :

```
FGT# diagnose sniffer packet port_name "ether host 01-00-0C-CC-CC-CC"
```

Configuration Example



Step 1: Create VLANs and forwarding domains

```

config system interface
    edit "vlan102_intern"
        set forward-domain 102
        set interface "port2"
        set vlanid 102
    next
    edit "vlan102_extern"
        set forward-domain 102
        set interface "port3"
        set vlanid 102
    next
    edit "vlan103_intern"
        set forward-domain 103
        set interface "port2"
        set vlanid 103
    next
    edit "vlan103_extern"
        set forward-domain 103
        set interface "port3"
        set vlanid 103
    next
end

```

Step 2: Create the appropriate Firewall Policies

```
config firewall policy
```

```
edit 1
    set srcintf "vlan102_extern"
    set dstintf "vlan102_intern"
        set srcaddr "all"
        set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY" next
edit 2
    set srcintf "vlan102_intern"
    set dstintf "vlan102_extern"
        set srcaddr "all"
        set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY" next
edit 3
    set srcintf "vlan103_intern"
    set dstintf "vlan103_extern"
        set srcaddr "all"
        set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY" next
edit 4
    set srcintf "vlan103_extern"
    set dstintf "vlan103_intern"
        set srcaddr "all"
        set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY" next
end
```

Firewalls and Security in Transparent Mode

This section contains information about using firewalls and security scanning in Transparent mode. It contains the following topics:

- [Firewall Policy Look Up](#)
- [Firewall Session List](#)
- [Security Scanning](#)
- [Integrated Switch Fabric Access Control List Policies](#)

Firewall Policy Look Up

In Transparent mode, like in NAT/Route mode, a firewall policy look up is based on the source and destination interfaces. The matching firewall policy will tell which actions to apply to the traffic, including logging and security scanning.

The FortiGate proceeds as follows to look for a matching firewall policy in Transparent mode:

- **Step 1:** an Ethernet IP frame ingresses a port (or a VLAN on a port), corresponding to a specific bridge instance (from the port VDOM and Forwarding domain). This frame contains a destination MAC address that we will call MAC_D.
- **Step 2:** The FortiGate is making a MAC_D address lookup in the bridge instance to determine the port where MAC_D has been learned. This will be the destination interface.
- **Step 3:** The FortiGate is then looking for a firewall policy corresponding to the couple < source interface + destination interface >. If multiple policies with the same couple < source interface + destination interface > exist, the FortiGate screens all of them from TOP to BOTTOM (as displayed in the configuration), until a match is found. It is important to make sure that the most specific firewall policies are located at the top of the policy list, to make sure that traffic is matched to the appropriate policy.

Firewall Session List

The flag **br** in the state line will indicate that this is a “bridged” session. See example below :

```
FGT# diagnose sys session list

session info: proto=17 proto_state=00 duration=59 expire=128 timeout=0 flags=000
00000 sockflag=00000000 sockport=0 av_idx=0 use=4
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 hakey=0
policy_dir=0 tunnel=/
state=may_dirty br rem
statistic(bytes/packets/allow_err): org=385/5/1 reply=0/0/0 tuples=2
orgin->sink: org pre->post, reply pre->post dev=3->4/4->3 gwy=0.0.0.0/0.0.0.0
hook=pre dir=org act=noop 192.168.182.93:1025->4.2.2.1:53(0.0.0.0:0)
hook=post dir=reply act=noop 4.2.2.1:53->192.168.182.93:1025(0.0.0.0:0)
```

```
misc=0 policy_id=1 auth_info=0 chk_client_info=0 vd=1 serial=000006d3 tos=ff/ff
imp2p=0 app=0
dd_type=0 dd_rule_id=0
```

Security Scanning

Security scanning occurs in the same manner in NAT/Route mode and Transparent mode. When a protection profile is enabled on a firewall policy for content inspection, the FortiGate acts like a transparent proxy for the protocols that need to be inspected.

The FortiGate will therefore intercept the TCP sessions and create its own session from client to server and server to client. The source and destination MAC addresses of the original L2 frames are however not altered in this communication, as described in the section Network operation : source MAC addresses in frames sent by or through the FortiGate.



Devices in the network communicating through the FortiGate do not know the presence of the FortiGate.

For more information about security scanning, see the Security Profiles handbook.

Integrated Switch Fabric Access Control List Policies



Traffic accepted and forwarded by an ISF policy is not subject to security inspection. Normally, you should only create ISF policies for traffic that you consider very low risk.

On FortiGate models that include NP4 and XLR ports and an integrated switch fabric, you can create an integrated switch fabric (ISF) access control list (ACL) firewall policy that allows some traffic (for example, multicast traffic) to bypass security inspection, resulting in reduced CPN and NP4 processor load.

This feature is only available in Transparent mode and only between port pairs.

Use the following command to add an ISF ACL shortcut policy:

```
config firewall isf-acl
  config port-pair-1
    edit 1
      set type binary
      set ingressport {port1 | port2}
      set offset
      set length
      set matchpattern <pattern in hex>
      set action {bypass | block}
    next
  edit 2
    set type 5-tuple
    set srcaddr: a.b.c.d/32
    set dstaddr 239.A.A.a/32
```

```
        set proto UDP
        set port XXX
        set action {bypass | block}
    end
end
```

IPsec VPN in Transparent Mode

This section contains information about configuring IPsec virtual private networks (VPNs) in Transparent mode. It contains the following topics:

- [Using IPsec VPNs in Transparent Mode](#)
- [Example 1: Remote Sites with Different Subnets](#)
- [Example 2: Remote Sites on the Same Subnet](#)

Using IPsec VPNs in Transparent Mode

In Transparent mode, IPsec VPN is supported in Policy-based configuration mode only.

IPsec VPN in Transparent mode can be used in those scenarios:

- Encrypt data over routed networks without changing anything on the routers. See example 1.
- Encrypt data over a non-routed transport network (extension of a LAN for example). See example 2.

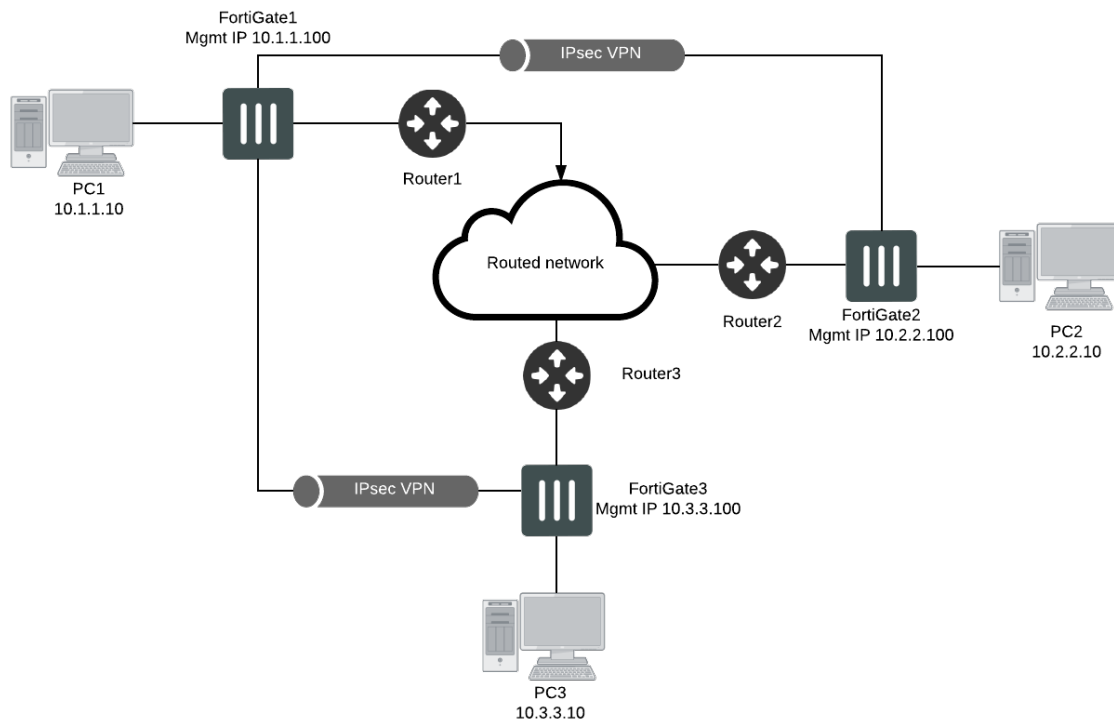
The following rules apply to IPsec in Transparent mode:

- If both remote FortiGate IPsec gateways are not in the same broadcast domain (separated by routers):
 - The hosts on each side must be on different subnets.
 - The FortiGate management IP addresses must be in the same subnet as the local hosts. This is the preferred option.
- If both remote FortiGate IPsec gateways are in the same broadcast domain (separated by optical switches for examples), the hosts on each side can be :
 - On the same subnet
 - On different subnet if the appropriate static route is configured on the remote Fortigate
 - The FortiGate management IP addresses can be in any different subnet than the local hosts
- A firewall Policy with the action IPsec is used to send traffic to the remote device into the tunnel.
Therefore, it is important to place all remote devices on the appropriate ports of the Fortigate to allow a proper match < source interface + destination interface > . See section Transparent mode Firewall processing for more details.



This scenario requires that the remote hosts located on the remote FortiGate's protected subnets have their MAC addresses hardcoded in FortiGate's static MAC entry list. If this is not configured then it is expected to see outage in network communications.

Example 1: Remote Sites with Different Subnets



This example provides a configuration example for IPsec VPN tunnels between three FortiGate in Transparent Mode in different subnets, as well as some troubleshooting steps.

The expectation for this example is that PC1 will be able to communicate via the IPsec tunnels with PC2 and PC3, which are in different subnets.

The requirements for this example are:

- Because both FortiGate are not in the same broadcast domain (separated by routers), the hosts on each side must be on different subnets.
- FortiGate management IP addresses must be in the same subnet as the local hosts
- The default gateways (router1, router2, router3) for PC1, PC2 and PC3 must be behind port2 in order for the FortiGate to match the appropriate Encrypt firewall policy (port1 --> port2)

Configuration of FortiGate 1 (FGT1):

Only relevant parts of configuration are provided.

```
config system settings
  set opmode transparent
  set manageip 10.1.1.100/255.255.255.0
end

config router static
  edit 1
```

```
        set gateway 10.1.1.254
    next
end

config firewall address
    edit "10.1.1.0/24"
        set subnet 10.1.1.0 255.255.255.0
    next
    edit "10.2.2.0/24"
        set subnet 10.2.2.0 255.255.255.0
    next
    edit "10.3.3.0/24"
        set subnet 10.3.3.0 255.255.255.0
    next
end

config vpn ipsec phase1
    edit "to_FGT2"
        set proposal 3des-sha1 aes128-sha1 des-md5
        set remote-gw 10.2.2.100
        set psksecret fortinet
    next
    edit "to_FGT3"
        set proposal 3des-sha1 aes128-sha1 des-md5
        set remote-gw 10.3.3.100
        set psksecret fortinet
    next
end

config vpn ipsec phase2
    edit "to_FGT2"
        set keepalive enable
        set phasename "to_FGT2"
        set proposal 3des-sha1 aes128-sha1
        set dst-subnet 10.2.2.0 255.255.255.0
        set src-subnet 10.1.1.0 255.255.255.0
    next
    edit "to_FGT3"
        set keepalive enable
        set phasename "to_FGT3"
        set proposal 3des-sha1 aes128-sha1
        set dst-subnet 10.3.3.0 255.255.255.0
        set src-subnet 10.1.1.0 255.255.255.0
    next
end

config firewall policy
    edit 1
        set srcintf "port1"
        set dstintf "port2"
        set srcaddr "10.1.1.0/24"
        set dstaddr "10.2.2.0/24"
        set action ipsec
        set schedule "always"
        set service "ANY"
        set inbound enable
        set outbound enable
```



```
        set vpntunnel "to_FGT2"
    next
    edit 3
        set srcintf "port1"
        set dstintf "port2"
        set srcaddr "10.1.1.0/24"
        set dstaddr "10.3.3.0/24"
        set action ipsec
        set schedule "always"
        set service "ANY"
        set inbound enable
        set outbound enable
        set vpntunnel "to_FGT3"
    next
end
```

Configuration of FortiGate 2 (FGT2):

Only relevant parts of configuration are provided.

```
config system settings
    set opmode transparent
    set manageip 10.2.2.100/255.255.255.0
end

config router static
    edit 1
        set gateway 10.2.2.254
    next
end

config firewall address
    edit "10.1.1.0/24"
        set subnet 10.1.1.0 255.255.255.0
    next
    edit "10.2.2.0/24"
        set subnet 10.2.2.0 255.255.255.0
    next
end

config vpn ipsec phase1
    edit "to_FGT1"
        set nattraversal disable
        set proposal 3des-sha1 aes128-sha1 des-md5
        set remote-gw 10.1.1.100
        set psksecret fortinet
    next
end

config vpn ipsec phase2
    edit "to_FGT1"
        set keepalive enable
        set phase1name "to_FGT1"
        set proposal 3des-sha1 aes128-sha1
        set dst-subnet 10.1.1.0 255.255.255.0
        set src-subnet 10.2.2.0 255.255.255.0
    next
```

```

end

config firewall policy
  edit 1
    set srcintf "port1"
    set dstintf "port2"
    set srcaddr "10.2.2.0/24"
    set dstaddr "10.1.1.0/24"
    set action ipsec
    set schedule "always"
    set service "ANY"
    set inbound enable
    set outbound enable
    set vpntunnel "to_FGT1"
  next
end

```

Troubleshooting procedure

All steps given when PC1 pings PC2.

Verify if IPsec tunnels are up

```

FGT1 # diagnose vpn tunnel list

list all ipsec tunnel in vd 0
-----
name=to_FGT2 ver=0 serial=1 10.1.1.100:0->10.2.2.100:0 lgwy=dyn tun=tunnel mode= auto
bound_if=0
proxyid_num=1 child_num=0 refcnt=7 ilast=0 olast=0
stat: rxp=0 txp=0 rxb=0 txb=0
dpd: mode=active on=1 idle=5000ms retry=3 count=0 seqno=1455
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=to_FGT2 proto=0 sa=0 ref=1 auto_negotiate=0 serial=5
  src: 10.1.1.0/255.255.255.0:0
  dst: 10.2.2.0/255.255.255.0:0

```

The above tunnel is down (output given as example)!

```

FGT2 # diagnose vpn tunnel list

list all ipsec tunnel in vd 0
-----
name=to_FGT1 10.2.2.100:0->10.1.1.100:0 lgwy=dyn tun=tunnel mode=auto bound_if=0
proxyid_num=1 child_num=0 refcnt=7 ilast=1 olast=1
stat: rxp=0 txp=0 rxb=0 txb=0
dpd: mode=active on=1 idle=5000ms retry=3 count=0 seqno=21 natt:
mode=none draft=0 interval=0 remote_port=0
proxyid=to_FGT1 proto=0 sa=1 ref=2 auto_negotiate=0 serial=1
  src: 10.2.2.0/255.255.255.0:0
  dst: 10.1.1.0/255.255.255.0:0
SA: ref=3 options=00000009 type=00 soft=0 mtu=1436 expire=1771 replaywin=0 seq no=1
life: type=01 bytes=0/0 timeout=1773/1800
dec: spi=cla8e951 esp=3des key=24 9213fdf22b150e01abb3535d1a647044eebf772b92f2f7ee
ah=sha1 key=20 66a38bf99f0b2d234f64b5a05187995c4f56f6bb
enc: spi=322067b4 esp=3des key=24 720e5680329937fb3630b7ed70bd41bb3114d3c269ae8b61

```

```
ah=sha1 key=20 e316113eb6ea03b014b3a5f9c1a3bd386637801a
```

The above tunnel is up!

Verify that destination local hosts are seen in the ARP table (necessary for IPsec despite being in TP mode)

```
FGT2 # get system arp
```

Address	Age(min)	Hardware Addr	Interface
10.2.2.10	2	00:50:56:00:76:04	root.b
10.2.2.254	0	00:09:0f:30:29:e4	root.b

Using the debug flow command on the initiator side (example on FortiGate1)

```
FGT1 # diagnose debug flow filter addr 10.1.1.10
FGT1 # diagnose debug flow show console enable
FGT1 # diagnose debug enable
FGT1 # diagnose debug flow trace start 50

FGT1 # id=36870 trace_id=615 msg="vd-root received a packet(proto=1, 10.1.1.10:512->10.2.2.10:8) from port1."
id=36870 trace_id=615 msg="allocate a new session-00000636"
id=36870 trace_id=615 msg="Allowed by Policy-1: encrypt"
id=36870 trace_id=615 msg="enter IPsec tunnel-to_FGT2"
id=36870 trace_id=615 msg="SA is not ready yet, drop"
id=36870 trace_id=616 msg="vd-root received a packet(proto=1, 10.1.1.10:512->10.2.2.10:8) from port1."
id=36870 trace_id=616 msg="Find an existing session, id-00000636, original direction"
id=36870 trace_id=616 msg="enter IPsec tunnel-to_FGT2"
id=36870 trace_id=616 msg="encrypted, and send to 10.2.2.100 with source 10.1.1.100"
id=36870 trace_id=616 msg="send out via dev-port2, dst-mac-00:09:0f:30:29:e0"
id=36870 trace_id=617 msg="vd-root received a packet(proto=1, 10.1.1.10:512->10.2.2.10:8) from port1."
id=36870 trace_id=617 msg="Find an existing session, id-00000636, original direction"
id=36870 trace_id=617 msg="enter IPsec tunnel-to_FGT2"
id=36870 trace_id=617 msg="encrypted, and send to 10.2.2.100 with source 10.1.1.100"
id=36870 trace_id=617 msg="send out via dev-port2, dst-mac-00:09:0f:30:29:e0"
id=36870 trace_id=618 msg="vd-root received a packet(proto=1, 10.2.2.10:512->10.1.1.10:0) from port2."
id=36870 trace_id=618 msg="Find an existing session, id-00000636, reply direction"
id=36870 trace_id=618 msg="send out via dev-port1, dst-mac-00:50:56:00:76:03"
id=36870 trace_id=619 msg="vd-root received a packet(proto=1, 10.1.1.10:512->10.2.2.10:8) from port1."
id=36870 trace_id=619 msg="Find an existing session, id-00000636, original direction"
id=36870 trace_id=619 msg="enter IPsec tunnel-to_FGT2"
id=36870 trace_id=619 msg="encrypted, and send to 10.2.2.100 with source 10.1.1.100"
id=36870 trace_id=619 msg="send out via dev-port2, dst-mac-00:09:0f:30:29:e0"
id=36870 trace_id=620 msg="vd-root received a packet(proto=1, 10.2.2.10:512->10.1.1.10:0) from port2."
id=36870 trace_id=620 msg="Find an existing session, id-00000636, reply direction"
id=36870 trace_id=620 msg="send out via dev-port1, dst-mac-00:50:56:00:76:03"
```



The message "id=36870 trace_id=615 msg="SA is not ready yet, drop" simply means that the tunnel was not up yet.

Using the debug flow command on the receiver side (example on FortiGate2)

```
FGT2 # diagnose debug flow filter addr 10.1.1.10
FGT2 # diagnose debug flow show console enable
FGT2 # diagnose debug enable
FGT2 # diagnose debug flow trace start 50

FGT2 # id=36870 trace_id=51 msg="vd-root received a packet(proto=1, 10.1.1.10:512->10.2.2.10:8) from port2."
id=36870 trace_id=51 msg="allocate a new session-00000435"
id=36870 trace_id=51 msg="Allowed by Policy-1:"
id=36870 trace_id=51 msg="send out via dev-port1, dst-mac-00:50:56:00:76:04"
id=36870 trace_id=52 msg="vd-root received a packet(proto=1, 10.2.2.10:512->10.1.1.10:0) from port1."
id=36870 trace_id=52 msg="Find an existing session, id-00000435, reply direction"
id=36870 trace_id=52 msg="enter IPsec tunnel-to_FGT1"
id=36870 trace_id=52 msg="encrypted, and send to 10.1.1.100 with source 10.2.2.100"
id=36870 trace_id=52 msg="send out via dev-port2, dst-mac-00:09:0f:30:29:e4"
```

Using the sniffer trace (example on FortiGate2)

```
FGT2 # diagnose sniffer packet any "host 10.2.2.10" 4

9.460021 root.b out arp who-has 10.2.2.10 tell 10.2.2.100
9.460028 port2 out arp who-has 10.2.2.10 tell 10.2.2.100
9.460034 port1 out arp who-has 10.2.2.10 tell 10.2.2.100
9.460462 port1 in arp reply 10.2.2.10 is-at 0:50:56:0:76:4
9.460462 root.b in arp reply 10.2.2.10 is-at 0:50:56:0:76:4
[...]
49.477368 port2 in 10.1.1.10 -> 10.2.2.10: icmp: echo request
49.477444 port1 out 10.1.1.10 -> 10.2.2.10: icmp: echo request
49.477898 port1 in 10.2.2.10 -> 10.1.1.10: icmp: echo reply
50.510023 port2 in 10.1.1.10 -> 10.2.2.10: icmp: echo request
50.510079 port1 out 10.1.1.10 -> 10.2.2.10: icmp: echo request
50.510524 port1 in 10.2.2.10 -> 10.1.1.10: icmp: echo reply
```



The above ARP process in Transparent mode with IPsec is allowing the Fortigate to:

- Identify the MAC address of the destination device 10.2.2.10
- Populate the MAC table (see below), which in turn will give a destination interface and allow a Firewall policy look-up


Check the FDB entries for the destination

```
FGT2 # diagnose netlink brctl name host root.b

port no   device  devname  mac addr          ttl
[...]
```

port no	device	devname	mac addr	ttl
1	2	port1	00:50:56:00:76:04	0

Example 2: Remote Sites on the Same Subnet



This example provides a configuration example for IPsec VPN tunnels between two FortiGate in Transparent Mode in the same subnet separated by a L2 transparent network and one remote subnet on the second site.



This scenario requires that PC1's MAC address is added to the FortiGate's static MAC table. The preferred scenario would be to have a router installed between the 2 FortiGate's.

The expectation for this example is that PC1 will be able to communicate via the IPsec tunnel with Server1 in the same subnet, and Server2 in a different subnet.

The requirements for this example are:

- The default gateway (FGT3) for PC1 and all remote device must be behind port2 of FGT1, in order for this FortiGate to match the appropriate Encrypt firewall policy (port1 --> port2)
- Despite being in Transparent mode, **FGT2 must have a valid route to Server2**
- FGT3 is used as a router between subnet 10.1.1.0/24 and 10.3.3.0/24.

PC1 MAC address added to FGT2 static MAC entries.

Server1 MAC address added to FGT1 static MAC entries.

Configuration of FortiGate 1 (FGT1):

Only relevant parts of configuration are provided.

```
config system settings
  set opmode transparent
  set manageip 10.1.1.100/255.255.255.0
end

config router static
  edit 1
    set gateway 10.1.1.252
  next
```

```
end

config system mac-address-table
    edit 00:50:56:00:76:04 ==>Server1
        set interface port2
    next
end

config firewall address
    edit "all"
    next
    edit "Server1"
        set subnet 10.1.1.20 255.255.255.255
    next
    edit "Server2"
        set subnet 10.3.3.30 255.255.255.255
    next
    edit "10.1.1.0/24"
        set subnet 10.1.1.0 255.255.255.0
    next
    edit "gateway"
        set subnet 10.1.1.254 255.255.255.255
    next
end

config vpn ipsec phase1
    edit "to_FGT2"
        set proposal 3des-sha1 aes128-sha1 des-md5
        set remote-gw 10.1.1.200
        set psksecret fortinet
    next
end

config vpn ipsec phase2
    edit "to_FGT2"
        set keepalive enable
        set phaselname "to_FGT2"
        set proposal 3des-sha1 aes128-sha1
        set src-subnet 10.1.1.0 255.255.255.0
    next
end

config firewall policy
    edit 1
        set srcintf "port1"
        set dstintf "port2"
        set srcaddr "10.1.1.0/24"
        set dstaddr "Server1"
        set action ipsec
        set schedule "always"
        set service "ANY"
        set inbound enable
        set outbound enable
        set vpntunnel "to_FGT2"
    next
    edit 2
        set srcintf "port1"
```

```
    set dstintf "port2"
    set srcaddr "10.1.1.0/24"
    set dstaddr "Server2"
    set action ipsec
    set schedule "always"
    set service "ANY"
    set inbound enable
    set outbound enable
    set vpntunnel "to_FGT2"
next
edit 3
    set srcintf "port1"
    set dstintf "port2"
    set srcaddr "10.1.1.0/24"
    set dstaddr "gateway"
    set action ipsec
    set schedule "always"
    set service "ANY"
    set inbound enable
    set outbound enable
    set vpntunnel "to_FGT2"
next
end
```



Firewall Policy 3 is not mandatory and is only used to allow PC1 to test a ping reachability to its default gateway 10.1.1.254.

Configuration of FortiGate 2 (FGT2):

Only relevant parts of configuration are provided.

```
config system settings
    set opmode transparent
    set manageip 10.1.1.200/255.255.255.0
end

config router static
    edit 1
        set gateway 10.1.1.252
    next
    edit 2
        set dst 10.3.3.0 255.255.255.0
        set gateway 10.1.1.254
    next
end

config system mac-address-table
    edit 00:50:56:00:76:03
        set interface wan1
    next
end

config firewall address
```

```
edit "all"
next
edit "PC1"
    set subnet 10.1.1.10 255.255.255.255
next
edit "10.1.1.0/24"
    set subnet 10.1.1.0 255.255.255.0
next
edit "10.3.3.0/24"
    set subnet 10.3.3.0 255.255.255.0
next
end

config vpn ipsec phase1
    edit "to_FGT1"
        set proposal 3des-sha1 aes128-sha1 des-md5
        set remote-gw 10.1.1.100
        set psksecret fortinet
    next
end

config vpn ipsec phase2
    edit "to_FGT1"
        set keepalive enable
        set phasename "to_FGT1"
        set proposal 3des-sha1 aes128-sha1
        set dst-subnet 10.1.1.0 255.255.255.0
    next
end

config firewall policy
    edit 1
        set srcintf "internal"
        set dstintf "wan1"
        set srcaddr "10.1.1.0/24" set dstaddr "PC1"
        set action ipsec
        set schedule "always"
        set service "ANY" set inbound enable
        set outbound enable
        set vpntunnel "to_FGT1"
    next
    edit 2
        set srcintf "internal"
        set dstintf "wan1"
        set srcaddr "10.3.3.0/24" set dstaddr "PC1"
        set action ipsec
        set schedule "always"
        set service "ANY" set inbound enable
        set outbound enable
        set vpntunnel "to_FGT1"
    next
end
```


Troubleshooting procedure

Check the ARP entries of PC1

```
C:\ arp -a
```

```
Interface: 10.1.1.10 --- 0x20003
Internet Address      Physical Address      Type
10.1.1.20             00-50-56-00-76-04    dynamic
10.1.1.254            00-09-0f-85-3f-c8    dynamic
```



MAC address **00-09-0f-85-3f-c8** is the FGT3 interface in subnet 10.1.1.0/24.

FDB entries of FGT1

```
FGT1 (global) # diagnose netlink brctl name host Vdom1.b
```

```
show bridge control interface Vdom1.b host. fdb:
size=256, used=6, num=6, depth=1
```

```
Bridge Vdom1.b host table
```

port no	device	devname	mac addr	ttl	attributes
1	10	port1	00:50:56:00:76:03	0	
2	9	port2	00:50:56:00:76:04	44	static
2	9	port2	00:09:0f:85:3f:c8	13	
1	10	port1	00:09:0f:88:2f:69	0	Local Static
2	9	port2	00:09:0f:88:2f:68	0	Local Static
2	9	port2	00:09:0f:23:01:d6	0	



MAC address **00:09:0f:23:01:d6** is “internal” port MAC address of FGT2 00:09:0F:23:01:D6. This is the MAC address used for management in the Transparent mode VDOM of FGT2, chosen between the lowest MAC address between wan1 (00:09:0F:78:00:74) and internal (00:09:0F:23:01:D6).

ARP entries of FGT2

```
FGT2 (TP) # get system arp
```

Address	Age (min)	Hardware Addr	Interface
10.1.1.20	82	00:50:56:00:76:04	TP.b
10.1.1.100	13	00:09:0f:88:2f:68	TP.b
10.1.1.254	76	00:09:0f:85:3f:c8	TP.b



it is important to have the entry for 10.1.1.254 which is the route to 10.3.3.0/24 .

IPsec Tunnel verification on FGT1

```

FGT1 (Vdom1) # diagnose vpn tunnel list

list all ipsec tunnel in vd 3
-----
name=to_FGT2 10.1.1.100:0->10.1.1.200:0 lgwy=dyn tun=tunnel mode=auto bound_if=0
proxyid_num=1 child_num=0 refcnt=10 ilast=0 olast=0
stat: rxp=2754 txp=2945 rxb=308448 txb=176700
dpd: mode=active on=1 idle=5000ms retry=3 count=0 seqno=166 natt:
mode=none draft=0 interval=0 remote_port=0
proxyid=to_FGT2 proto=0 sa=1 ref=2 auto_negotiate=0 serial=1
src: 10.1.1.0/255.255.255.0:0
dst: 0.0.0.0/0.0.0.0:0
SA: ref=3 options=00000009 type=00 soft=0 mtu=1436 expire=1271 replaywin=0 seqno=1e1
life:type=01 bytes=0/0 timeout=1750/1800
dec: spi=3f148cb7 esp=3des key=24 834832201a0dbbf60b0098106f08380538dbd94cacdlad31
ah=sha1 key=20 b0257a135cba745b956bef3d4b8a6e65934c074b
enc: spi=1895305e esp=3des key=24 4d3092f0b3f84184d4779f85a9953230bf9bc28bd93c0afa
ah=sha1 key=20 0c70acf6ad2193ec5934e2a4332fd09f32016e60
npu_flag=00 npu_rgw=10.1.1.200 npu_lgwy=10.1.1.100 npu_selid=0

```

Sniffer trace on FGT1 when PC1 pings all 3 remote destinations

```

FGT1 (Vdom1) # diagnose sniffer packet any "icmp" 4

interfaces=[any]
filters=[icmp]
0.342268 port1 in 10.1.1.10 -> 10.3.3.30: icmp: echo request
0.342844 port2 in 10.3.3.30 -> 10.1.1.10: icmp: echo reply
0.342884 port1 out 10.3.3.30 -> 10.1.1.10: icmp: echo reply
0.771700 port1 in 10.1.1.10 -> 10.1.1.20: icmp: echo request
0.772504 port2 in 10.1.1.20 -> 10.1.1.10: icmp: echo reply
0.772539 port1 out 10.1.1.20 -> 10.1.1.10: icmp: echo reply
0.907377 port1 in 10.1.1.10 -> 10.1.1.254: icmp: echo request
0.907850 port2 in 10.1.1.254 -> 10.1.1.10: icmp: echo reply
0.907883 port1 out 10.1.1.254 -> 10.1.1.10: icmp: echo reply

```

Sniffer trace on FGT1 filtered on IPsec protocol

```

FGT1 (Vdom1) # diagnose sniffer packet port2 "proto 50" 6

interfaces=[port2]
filters=[proto 50]
pcap_lookupnet: port2: no IPv4 address assigned

1.249003 port2 -- 10.1.1.100 -> 10.1.1.200: ip-proto-50 92
0x0000 0009 0f23 01d6 0009 0f88 2f68 0800 4500 ...# ..... /h..E.
0x0010 0070 c9e6 0000 3f32 9a48 0a01 0164 0a01 .p ... ?2.H...d..
0x0020 01c8 1895 305f 0000 01e2 02b6 37b6 8b2c ....0_ ..... 7...,

1.249478 port2 -- 10.1.1.200 -> 10.1.1.100: ip-proto-50 92
0x0000 0009 0f88 2f68 0009 0f23 01d6 0800 4500 ....h...# ... E.
0x0010 0070 2e31 0000 3f32 35fe 0a01 01c8 0a01 .p.1..?25 .....
0x0020 0164 3f14 8cb8 0000 01e2 324d 66e2 9236 .d? ..... 2Mf..6

```



From the above trace, the MAC address **0009 0f88 2f68** is the MAC address of FGT1 port2 . This is the MAC address used for management in the Transparent mode VDOM of FGT1, chosen between the lowest MAC address between port1 (00:09:0F:88:2F:69) and port2 (00:09:0F:88:2F:68).

Debug flow on FGT1 filtered on Server3

```
FGT1 (Vdom1) # diagnose debug flow filter addr 10.3.3.30
FGT1 (Vdom1) # diagnose debug flow show console enable
FGT1 (Vdom1) # diagnose debug enable
FGT1 (Vdom1) # diagnose debug flow trace start 10

id=20085 trace_id=11 msg="vd-Vdom1 received a packet(proto=1, 10.1.1.10:512->10.3.3.30:8)
from port1."
id=20085 trace_id=11 msg="Find an existing session, id-00004e85, original direction"
id=20085 trace_id=11 msg="enter IPsec tunnel-to_FGT2"
id=20085 trace_id=11 msg="encrypted, and send to 10.1.1.200 with source 10.1.1.100"
id=20085 trace_id=11 msg="send out via dev-port2, dst-mac-00:09:0f:23:01:d6"
id=20085 trace_id=12 msg="vd-Vdom1 received a packet(proto=1, 10.3.3.30:512->10.1.1.10:0)
from port2."
id=20085 trace_id=12 msg="Find an existing session, id-00004e85, reply direction"
id=20085 trace_id=12 msg="send out via dev-port1, dst-mac-00:50:56:00:76:03"
```



From the trace above, **dst-mac-00:09:0f:23:01:d6** is "internal" port MAC address of FGT2 00:09:0F:23:01:D6. This is the MAC address used for management in the Transparent mode VDOM of FGT2, chosen between the lowest MAC address between wan1 (00:09:0F:78:00:74) and internal (00:09:0F:23:01:D6).

Using FortiManager and FortiAnalyzer

FortiManager and FortiAnalyzer are supported similarly to NAT/Route mode. For more information about this please consult the Fortinet documentation at <http://docs.fortinet.com> or the Knowledge base at <http://kb.fortinet.com>.

Establishing a communication to a FortiAnalyzer is done as per the example hereafter (from global level if VDOM is enabled). This setting is independent from being in Transparent mode. However, as stated earlier in this section, the management VDOM must have IP connectivity to the FortiAnalyzer.

```
FGT (global) # show system fortianalyzer

config system fortianalyzer
    set status enable
    set server 10.2.2.2
end
```

High Availability in Transparent Mode

This section contains information about configuring high availability in Transparent mode. It contains the following topics:

- [Virtual Clustering](#)
- [MAC Address Assignment](#)



For complete information about HA, please refer to the Fortigate Administration Guide or the HA Technical guides available at <http://docs.fortinet.com> or the Knowledge Base at <http://kb.fortinet.com>.

Any other statement and feature description in this document apply to a FortiGate Cluster running in Active-Passive mode.

Virtual Clustering

If VDOM (virtual domain) is enabled on a cluster operating Transparent Mode, HA Virtual Clustering can be configured in active-passive mode.

This will provide:

- Failover protection between two instances of a VDOM operating on two different FortiGate in the cluster.
- Load balancing between the FortiGate units on a per-VDOM basis.

The roles have been defined such as, in normal operation:

- FortiGate1 is Master for Vdom1 and Slave for Vdom2
- FortiGate2 is Master for Vdom2 and Slave for Vdom1

In case of a failure or reboot of a FortiGate, the remaining unit will become Master for Vdom1 and Vdom2.



The VDOMs given in this example are showing physical ports but a VDOM can also include VLAN interfaces.



The L2 connectivity between the FortiGate is showing 4 separate L2 switches, but it could also be one single switch on each side configured with appropriate VLANs.

Configuration example

- FortiGate1:

```
FGT1 (global) # show system ha

config system ha
  set mode a-p
```

```
set hbdev "port5" 0 "port6" 0
set vcluster2 enable
set override disable
set priority 200
  config secondary-vcluster
    set override enable
    set priority 100
    set vdom "Vdom2"
  end
end
```

- **FortiGate2:**

```
FGT2 (global) # show system ha

config system ha
  set mode a-p
  set hbdev "port5" 0 "port6" 0
  set vcluster2 enable
  set override disable
  set priority 200
    config secondary-vcluster
      set override enable
      set priority 100
      set vdom "Vdom2"
    end
end
```

MAC Address Assignment

If a cluster is operating in Transparent mode, the FortiGate Clustering Protocol (FGCP) assigns a virtual MAC address for the Master unit management IP address. Since you can connect to the management IP address from any interface, all of the FortiGate interfaces appear to have the same virtual MAC address.

Best Practices

1. Create forwarding domains when VLANs are used and set `vlanforward to disable` on all relevant physical interface.
2. The forward-domain ID can be different to the VLAN ID, but it is recommended for troubleshooting and readability to keep them the same.
3. Only interfaces from the same forwarding domains can have firewall policies between each others.
4. In order to allow IVL (independent VLAN learning), the VLANs must be placed in separate forwarding domains.
5. If an out-of-band management is desired, use if possible a VDOM in NAT/Route mode as management VDOM and create (an) other Transparent mode VDOM(s) for the user traffic.
6. As Spanning Tree BPDUs are not forwarded by default, insert the FortiGate with caution to avoid L2 loops.
7. Multicast packets are not forwarded by default; this might cause routing protocols (RIP2, OSPF) disruption.
8. When using HSRP or VRRP configure static MAC entries for the Virtual MAC addresses.



Copyright© 2017 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features, or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.